

Service composition via simulation

Seminari di Ingegneria del SW
Slides by Fabio Patrizi and Giuseppe De Giacomo
DIS, Sapienza – Università di Roma

Rome - June, 2008

Essential overview

- Computing composition via simulation
- Using an LTL synthesis tool, TLV, for computing composition via simulation

The Problem

Given:

- a community of available services

$$\mathcal{C} = \{S_1, \dots, S_n\};$$

- a target service

$$T;$$

Find a *composition* (or *orchestrator*) s.t.

$$\mathcal{C} \text{ mimics } T$$

The Problem (cont.)

We model services as transition systems:

- A TS is a tuple $T = \langle A, S, s_0, \delta, F \rangle$ where:
 - A is the set of actions
 - S is the set of states
 - $s_0 \in S$ is the set of initial states
 - $\delta \subseteq S \times A \times S$ is the transition relation
 - $F \subseteq S$ is the set of final states

Finding a composition

Strategies for computing compositions:

- Reduction to PDL
- Simulation-based 

Simulation Relation

Intuition:

a service \mathcal{T} can be simulated by community \mathcal{C} if \mathcal{C} can reproduce \mathcal{T} 's behavior over time.

Simulation Relation (cont.)

- Given two transition systems $\mathcal{T} = \langle A, T, t^0, \delta_T, F_T \rangle$ and $\mathcal{C} = \langle A, S, s_C^0, \delta_C, F_C \rangle$ a **simulation** relation on $T \times C$ is a binary relation on the states $t \in T$ and s of C such that:
 - $(t, s) \in R$ implies that
 - t is *final* implies that s is *final*
 - for all actions a
if $t \xrightarrow{a} t'$ then $\exists s' . s \xrightarrow{a} s'$ and $(t', s') \in R$
 - If **exists a simulation** relation R (such that $(t^0, s_C^0) \in R$, then we say that or **T is simulated by C** (or **C simulates T**).
 - **Simulated by** is (i) a simulation; (ii) the largest simulation
- *NB1: Simulated by is a co-inductive definition!*
- *NB2: A simulation is just one of the two directions of a bisimulation*

Simulation Relation (cont.)

Algorithm ComputingSimulation

Input: transition system $T = \langle A, T, t^0, \delta_T, F_T \rangle$ and transition system $C = \langle A, S, s_C^0, \delta_C, F_C \rangle$

Output: the **simulated-by** relation (the largest simulation)

Body

$R = \emptyset$

$R' = T \times S - \{(t, s) \mid t \in F_T \wedge \neg(s \in F_C)\}$

while ($R \neq R'$) {

$R := R'$

$R' := R' - \{(t, s) \mid \exists t', a. t \xrightarrow{a} t' \wedge \neg \exists s' . s \xrightarrow{a} s' \wedge (t', s') \in R'\}$

}


return R'

Ydob


■

Simulation relation (cont.)

T:



C:



Can \mathcal{C} simulate T?

YES!

Computing composition via simulation

Idea:

A service community can be seen as the (possibly N-DET) *asynchronous product* of available services...

Computing composition via simulation

Let S_1, \dots, S_n be the TSs of the component services.

The **Community TS** $\mathcal{C} = \langle A, S_{\mathcal{C}}, s_{\mathcal{C}}^0, \delta_{\mathcal{C}}, F_{\mathcal{C}} \rangle$ is the **asynchronous product** of S_1, \dots, S_n where:




- A is the set of actions
- $S_{\mathcal{C}} = S_1 \times \dots \times S_n$
- $s_{\mathcal{C}}^0 = (s_1^0, \dots, s_n^0)$
- $F \subseteq F_1 \times \dots \times F_n$
- $\delta_{\mathcal{C}} \subseteq S_{\mathcal{C}} \times A \times S_{\mathcal{C}}$ is defined as follows:

$(s_1 \times \dots \times s_n) \xrightarrow{a} (s'_1 \times \dots \times s'_n)$ iff


- $\exists i. s_i \xrightarrow{a} s'_i \in \delta_i$
- $\forall j \neq i. s'_j = s_j$

Computing composition via simulation (cont.)

Available services



Community TS



Computing composition via simulation (cont.)

Idea:

Theorem:

A composition exists if and only if
 \mathcal{C} simulates T


including all behaviors of any feasible

... thus, the problem becomes:

“Can the community TS \mathcal{C} simulate
target service T?”

Computing composition via simulation (cont.)

Community TS



The orchestrator generator


- Given the largest simulation \mathbf{S} from TSt to TS_c (which include the initial states), we can build the **orchestrator generator**.
- This is an orchestrator program that can change its behavior reacting to the information acquired at run-time.

Def: **OG** = $\langle A, [1, \dots, n], S_r, S_r^0, \omega_r, F_r \rangle$ with

- A : the **actions** shared by the community
- $[1, \dots, n]$: the **identifiers** of the available services in the community
- $S_r = S_{t_1} \times \dots \times S_{t_n}$: the **states** of the orchestrator program
- $S_r^0 = (s_{t_1}^0, s_{t_2}^0, \dots, s_{t_n}^0)$: the **initial state** of the orchestrator program
- $F_r \subseteq \{ (s_{t_1}, s_{t_2}, \dots, s_{t_n}) \mid s_{t_i} \in F_{t_i} \}$: the **final states** of the orchestrator program
- $\omega_r: S_r \times A_r \rightarrow [1, \dots, n]$: the **service selection function**, defined as follows:
 - If $s_{t_i} \rightarrow_a s'_{t_i}$ then **choose** k s.t. $\exists s_k. s_k \rightarrow_a s'_k \wedge (s'_{t_i}, (s_1, \dots, s'_k, \dots, s_n)) \in \mathbf{S}$
- $\delta_r \subseteq S_r \times A_r \times [1, \dots, n] \rightarrow S_r$: the **state transition function**, defined as follows:
 - Let $\omega_r(s_{t_1}, s_{t_2}, \dots, s_{t_n}, a) = k$ then
 $(s_{t_1}, s_{t_2}, \dots, s_{t_n}) \rightarrow_{a,k} (s'_{t_1}, s'_{t_2}, \dots, s'_{t_n})$ where $s_k \rightarrow_a s'_k$


The orchestrator generator (cont.)

- From the maximal simulation, we can easily derive an **orchestrator generator**, e.g.:



The orchestrator generator (cont.)

From OG, one can select services to perform client actions.



Rome - June, 2008

Web service composition via simulation and TLV

17

Computing composition via simulation (cont.)

Summing up:

- Compute community TS \mathcal{C} ;
- Compute the maximal simulation of \mathcal{T} by \mathcal{C} ;
- - If simulation exists, compute OG;
 - else return “unrealizable”;
- Exploit OG for available service selection, even in a *just-in-time* fashion.

Rome - June, 2008

Web service composition via simulation and TLV

18

On-the-fly failure recovery with OG [KR08]

OG already solves:

- **Temporary freezing** of an available service k
 - Stop selecting k in OG until service k comes back!
- **Unexpected state change** of an available service
 - Recompute OG / simulated-by from new initial state ...
 - ... but OG / simulated-by independent from initial state!
 - Simply use old OG / simulated-by from the new state!!

Parsimonious failure recovery with OG [KR08]

Algorithm ComputingSimulation - parametrized version

Input: transition system $T = \langle A, T, t^0, \delta_T, F_T \rangle$ and

transition system $C = \langle A, S, sc^0, \delta_C, F_C \rangle$

relation R_{init} including then simulated-by

relation R_{sure} included then simulated-by

Output: the **simulated-by** relation (the largest simulation)

Body

$R = \emptyset$

$R' = R_{init} - \{(t, s) \mid t \in F_t \wedge \neg(s \in F_C)\}$

while ($R \neq R'$) {

$R := R'$

$R' := R' - \{(t, s) \mid \exists t', a. t \rightarrow_a t' \wedge \neg \exists s'. s \rightarrow_a s' \wedge (t', s') \in R' \cup R_{sure}\}$

}

 return $R' \cup R_{sure}$

Ydob

Parsimonious failure recovery with OG (cont.) [KR08]

Let $[1, \dots, n] = WUF$ be the available services.

Let $R = R_{WUF}$ be the **simulated-by** relation of target by services WUF .

Then consider the following relations [KR08]:

- $R_W \subseteq \pi_W(R_{WUF})$
 - $(\pi_W(R))$ is not a simulation of target by services W)
 - $\pi_W(R_{WUF})$ is the **projection on W** of a relation: easy to compute
- $R_W \times F \subseteq R_{WUF}$
 - $(R_W \times F)$ is a simulation of target by services WUF
 - $R_W \times F$ is the **cartesian product** of 2 relations (F is trivial): easy to compute

Parsimonious failure recovery with OG (cont.) [KR08]

When **services F die**

compute simulated-by R_W starting $\pi_W(R_{WUF})$!

If **dead services F come back**

compute simulated-by R_{WUF} starting $R_W \times F$!

Remember:

- $R_W \subseteq \pi_W(R_{WUF})$
 - $(\pi_W(R))$ is not a simulation of target by services W)
- $R_W \times F \subseteq R_{WUF}$
 - $(R_W \times F)$ is a simulation of target by services WUF)

Comments

- *Full observability* is crucial for OG to work properly. In fact, in order to propose services for action execution, state of each available service *needs* to be known.
- *Partial observability* possible through knowledge operator **[to be done]**
- Interesting extension: dealing with nondeterministic (devilish) available services (a slightly different notion of simulation is needed). **[KR08]**
- OG allows for failure tolerance! **[KR08]**

Tools for computing composition based on simulation

- Computing composition via simulation
- Use simulation computing tools for composition **[to be done]**
- Use LTL-based synthesis tools, like TLV, for indirectly computing composition via simulation **[Patrizi PhD08]** 


Composing services via TLV

The environment TLV (Temporal Logic Verifier) [Pnueli and Shahar, 1996] is a useful tool that can be used to

automatically compute the orchestrator generator,

given a problem instance.

Composing services via TLV (cont.)




Composing services via TLV (cont.)

We provide TLV a file written in (a flavour of) SMV, a language for specifying TSs.


- SMV specifications are typically composed of *modules, properly interconnected*;
- Intuitively, a module is a *sort of TS* which may share variables with other modules;
- A module may contain several submodules, properly synchronized;
- Module **main** is mandatory and contains all relevant modules, properly interconnected and synchronized.

Composing services via TLV (cont.)



Composing services via TLV (cont.)


We introduce SMV formalization by means of the following example, proceeding top-down:



Composing services via TLV (cont.)

- The application is structured as follows:
 - 1 module **main**
 - 1 module **Output**, representing OG service selection
 - 1 module **Input**, representing the (synchronous) interaction community-target
 - 1 module **mT1** representing the target service
 - 1 module **mSi** per available service

Module interconnections




Rome - June, 2008

Web service composition via simulation and TLV

31

The module main

- Instance independent
- Includes synchronous submodules In and Out.



Rome - June, 2008

Web service composition via simulation and TLV

32

The module Output

- Depends on number of available services. In this case: 2

```
MODULE Output
VAR
    index:0..2;
ASSIGN
    init(index) := 0;
    next(index) := 1..2;
```


Number of available services

Only for init

The module Output (cont.)

```
MODULE Output
VAR
    index:0..2;
ASSIGN
    init(index) := 0;
    next(index) := 1..2;
```

```
MODULE main
VAR
    In: system Input(Out.index);
    Out: system Output;
DEFINE
    good := !In.failure;
```



The module Input

Action alphabet +
special action
nil (used for init)

Target service

Available service 1

Available service 2

```
MODULE Input(index)
VAR
  action : {nil,search,display,return};
  T1 : mT1(action);
  S1 : mS1(index,action);
  S2 : mS2(index,action);
DEFINE
  failure := (S1.failure | S2.failure) |
    !(T1.final -> (S1.final & S2.final));
```

Fail if:

- S1 or S2 (... or SN) fail, OR
- T1 can be in a final state when S1 or S2 (... or SN) are not.

The target module mT1


■ Think of mT1 as an action producer

TS States

Init

Output relation
(non-deterministic,
in general)

```
MODULE mT1(act)
VAR
  loc : 0..1;
  ASSIGN
    init(loc) := 0;
    init(act) := nil;
    next(loc) := case
      esac;
    next(act) := case
      esac;
  DEFINE
```



Transition function
(deterministic, in general)

State 0 is final

The target module mT1 (cont.)

1. A statement of the form:

```
next(loc):=
  case
    case_1;
    ...
    case_n;
    TRUE : loc;
  esac;
```

is included for defining next loc value. Each `case_i` expression refers to a different pair $s, a \in S_t \times A_t$ such that $\delta_t(s, a)$ is defined (order does not matter) and assumes the form:

$loc = ind(s) \ \& \ act = a : \delta_t(s, a)$

2. A statement of the form:

```
next(act):=
  case
    case_0;
    case_1;
    ...
    case_n;
    TRUE : act;
  esac;
```

is included for defining next act assignment. Let $act : S_t \rightarrow 2^{A_t}$ be defined as $act(s) = \{a \in A_t \mid \exists s' \in S_t \text{ s.t. } s' = \delta_t(s, a)\}$. Then, `case_0` assumes the form:

$act = nil : act(s_0)$


For $i > 0$, each `case_i` expression refers to a different pair $s, a \in S_t \times A_t$ such that $act(\delta_t(s, a)) \neq \emptyset$ (order does not matter) and assumes the form:

$loc = ind(s) \ \& \ act = a : act(\delta_t(s, a))$


The target module mT1 (cont.)

```
MODULE mT1(act)
VAR
  loc : 0..1;
ASSIGN
  init(loc) := 1;
  init(act) := nil;
  next(loc) := 
    CASE
      loc = 0 & act = search : 1;
      loc = 1 & act = display : 0;
      TRUE : loc;
    esac;
  next(act) := 
    CASE
      act = nil : {search};
      loc = 0 & act = search : {display};
      loc = 1 & act = display : {search};
      TRUE : {act};
    esac;
  DEFINE
    final := (loc = 0);
  END

MODULE Input(index)
VAR
  action : {nil,search,display,return};
  T1 : mT1(action);
  S1 : mS1(index,action);
  S2 : mS2(index,action);
  DEFINE
    failure := (S1.failure | S2.failure) |
      !(T1.final & (S1.final & S2.final));
```



The available service module mS1




Rome - June, 2008

Web service composition via simulation and TLV

39

The available service module mS2



Rome - June, 2008

Web service composition via simulation and TLV

40

Putting things together

```
MODULE main
VAR
  In: system Input(Out.index);
  Out: system Output;
DEFINE
  good := !In.failure;
```

Never changes

```
MODULE Output
VAR
  index:0..2;
ASSIGN
  init(index) := 0;
  next(index) := 1..2;
```

Number of available services

Putting things together (cont.)

```
MODULE Input(index)
VAR
  action : {nil,search,display,return};
  T1 : mT1(action);
  S1 : mS1(index,action);
  S2 : mS2(index,action);
DEFINE
  failure := (S1.failure | S2.failure) |
  !(T1.final -> (S1.final & S2.final));
```

Whole shared action alphabet plus special action nil

Never changes

Index changes, add one module per available service

Index changes, add one conjunct/disjunct per available service

Putting things together (cont.)

```

MODULE mT1(act)
  VAR
    loc : 0..1;
  ASSIGN
    init(loc) := 0;
    init(act) := nil;
    next(loc) :=
      case
        loc = 0 & act = search : 1;
        loc = 1 & act = display : 0;
        TRUE : loc;
      esac;
    next(act) :=
      case
        act = nil : {search};
        loc = 0 & act = search : {display};
        loc = 1 & act = display : {search};
        TRUE : {act};
      esac;
  DEFINE
    final := (loc = 0);

```

Target service states

Never changes

Depends on service, see general rules.

List final states using either logical OR '||'
(e.g., (loc=0||loc=1||loc=3)) or set construction (e.g., (loc={0,1,3})).

Rome - June, 2008

Web service composition via simulation and TLV

43

Putting things together (cont.)

```

MODULE mS1(index,action)
  VAR
    loc : 0..1;
  ASSIGN
    init(loc) := 0;
    next(loc) :=
      case
        index != 1 : loc;
        loc=0 & action in {search} : {0,1};
        loc=1 & action in {display,return} : {0};
        TRUE : loc;
      esac;
  DEFINE
    failure :=
      index = 1 &
      (index = 0 & action in {search} ) ||
      (index = 1 & action in {display, return});
    final := (loc = 0);

```

Available service states

Never changes

Depends on service, see general rules.

Index changes. Same as module name

Rome - June, 2008

Web service composition via simulation and TLV


44

Putting things together (cont.)

```
MODULE mS2(index,action)
  DEFINE
    failure :=
      index = 2 & !(action in {display});
    final := TRUE;
```

Running the specification

Running TLV with our specification as input...



Running the specification (cont.)

That is, the following OG:

